WHITE PAPER

Artificial Intelligence for dentistry

FDI ARTIFICIAL INTELLIGENCE WORKING GROUP

Falk Schwendicke, Markus Blatz, Sergio Uribe, William Cheung, Mahesh Verma, Jina Linton, and Young Jun Kim

VISIT…

Table of contents

Executive Summary

Artificial intelligence (AI) is defined as machines performing intellectual tasks, which usually were assumed can only be performed by humans.

Based on an exploding amount of available digital data, significant advances in hardware performance, as well as progress in algorithmic and software approaches, the current capacities of AI are unprecedented. In medicine and specifically in dentistry, a growing number of AI applications are being developed, and the number of research studies published in the field grows exponentially at present.

FDI World Dental Federation (FDI) acknowledges AI as one of the key technologies for the dental profession for the future and how it may foster FDI's *Vision 2030* strategy for delivering optimal oral health for all.

The present FDI White Paper on AI for oral and dental care, public health, dental education and research has four objectives:

(1) To define AI and lay out the technology behind it for the professional community.

(2) To systematize and comprehensively display cases where AI can benefit oral and dental healthcare systems/delivery, education, and research.

(3) To identify areas where AI can facilitate FDI's *Vision 2030* strategy.

(4) To highlight potential risks of AI and the need for the community to engage and enforce standards, regulations, and best practices of AI.

This White Paper describes the application of AI in four domains of dentistry. First, applications of AI in individual patient care (i.e., in dental offices) are discussed, including image analysis, data synthesis and predictive dentistry, evidence-supported planning and conductive therapies, and patient communication and interaction. Second, the role of data analytics and AI in community and public health, including oral health surveillance and the provision of care to those who cannot access traditional dentist services is described. Third, the role of AI for workforce planning and monitoring, the chances it offers for diversifying the workforce and the impact it may have on dental education are displayed. In parallel, the need to increase data and digital literacy of professionals is highlighted, and the risks of automation and complacency are discussed. Fourth, the chances of AI for dental research, particularly for analyzing complex and big datasets, but also for enabling what is called "Precision" or "Personalized" dentistry are described.

Notably, the White Paper also highlights the needs to balance data protection and accessibility, and to revise conventional peer review approaches when dealing with AI and complex data analytics research.

In addition to showing which applications are available or conceivable, the current risks and limitations of AI are discussed, including bias and limited generalizability of applications, the limited accessibility and interoperability of the data underlying AI, as well as the problems when implementing and maintaining AI in clinical care.

This White Paper also describes how governance and agreed principles are needed to ensure that AI works to people's benefits.

Finally, this White Paper concludes with the implications for clinical practice and a call-to-action: To make sure that AI applications acquire maximal usefulness and generalizability, data protection concerns need to be carefully balanced against the potential benefits of AI for patients and society. Technological approaches supporting this balance should be explored more in depth; data harmonization for interoperability should be actively pursued; AI applications with demonstrated benefit for patients, providers or the healthcare system should be promoted and incentivized; evidence supporting AI applications for oral and dental care needs to be strengthened, including evidence supporting true clinical utility, applicability and implementation of these

applications; providers and healthcare decision-makers need to acquire basic literacy about data, digital applications and AI; the community should engage into the widespread activities around standardization of AI.

This White Paper will serve as source of information and a basis for discussion within the dental profession, but also as groundwork for further activities of FDI in this direction.

Section 1. Background and Objectives

Background of Artificial Intelligence

Data is increasingly considered as a central component for future societies to thrive. Data-driven analytics is now seen as the fourth paradigm to understand nature, besides the experimental and theoretical sciences and simulation (Bell et al. 2009). The excitement around data science is facilitated by an exploding amount of structured and unstructured available data, but also by the software and hardware tools at hand to leverage these data.

For many modern data-related applications, artificial intelligence (AI) is the technique to use the data; it is the instrument to analyze especially large amounts of data and put them to use. The term AI was coined in the mid-1950s, and the English Oxford Living Dictionary defines AI as "the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages" (Oxford Learners Dictionaries 2022).

Most AI applications involve machine learning, where machines repeatedly analyze datasets to identify ("learn") an inherent pattern in the data, such as detecting objects on images or predicting events from tabular data. A specific subfield of machine learning, employed to analyze particularly complex data such as images or speech, is deep learning, where prior- tested multi-layered networks of mathematical operations are used to allow the representation of features (LeCun et al. 2015). By now, AI encompasses a range of applications. These include computer vision (for example face recognition by smartphones or security cameras, autonomous driving), natural language processing (voice assistants or chatbots that can systematize and link both structured and unstructured text and similar tools), robotics, virtual reality and simulation systems (for example in manufacturing or surgery but also drug development), and decision-making support (for example in medical symptom checkers). AI is abundant in our everyday world and with exponential growth of data, hardware capacities, and new software architectures has led to astonishing dynamics in many of the described fields.

In healthcare and health sciences, the generation of data, its digitalization and mining, have started with a significant lag compared to other industries. This is rooted in several aspects, data protection being a major one, but also because health data is generated and stored in silos and health data interoperability is limited. The evolution of tested, wearable sensors accelerated the need for AI for health applications. Only recently, data-driven medical applications involving AI have reached clinical maturity while the exploding numbers of research studies indicate that medicine will be faced with a wealth of AI tools swamping the market soon. A similar development will be seen in oral and dental healthcare and sciences (Schwendicke and Krois 2021).

White Paper Objectives

The present FDI World Dental Federation (FDI) White Paper on AI for oral and dental care, public health, dental education and research has the following objectives:

1. To define AI and lay out the technology behind it for the professional community; this has already been covered above in brief.

2. Systematize and comprehensively display cases where AI can benefit oral and dental healthcare systems/delivery, education, and research.

3. Identify areas where AI can facilitate FDI's *Vision 2030* strategy for delivering optimal oral health for all.

4. Highlight potential risks of AI and the need for the community to engage and enforce standards, regulations, and best practices of AI.

We will cover these aims along with the various domains where we see that AI will impact dentistry. These include clinical care provided by dental professionals in dental offices, dental community and public health, dental education, and dental research (Figure 1). We will conclude this FDI White Paper with a section on challenges to allow AI to fulfil its promises and a call to action for the oral and dental health community.

Figure 1. Network of domains and subdomains of AI applications in oral and dental healthcare provided by dental professionals to individual patients (blue), as part of community and public dental health (steel), dental education and workforce management (turquoise), and oral and dental research (grey). Sources of bias (data) and further determinants of AI applications in oral and dental healthcare (formats, frameworks and infrastructure) are indicated in green.
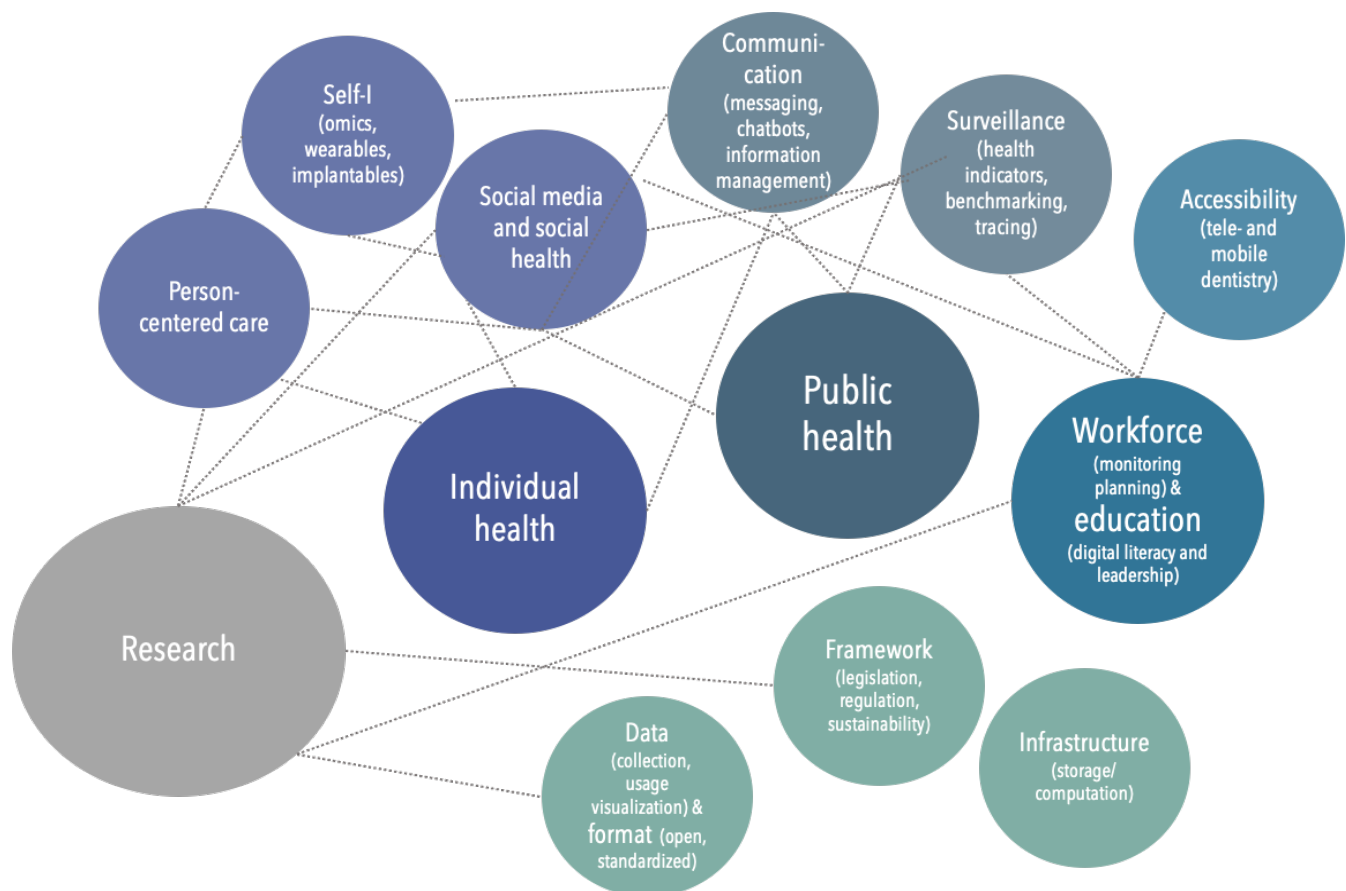
Section 2. Dentistry for individual patients: better understanding for better care

Helping clinicians provide better care on a consistent basis is one of the key goals of integrating AI in clinical dentistry. Data managing and AI applications are applied at all stages of the patient journey (Figure 2) and in all dental specialty areas. Four areas are to be highlighted; image analysis, data synthesis and prediction, evidence-supported treatment planning and conduct, and patient interaction, particularly during supportive care.

Image Analysis

Image interpretation in dental radiology is one of the currently most visible applications of AI in clinical dentistry. Most common pathologies such as caries or apical lesions, periodontal bone loss, cysts or bony fractures are detected by AI on 2D and, increasingly, 3D images. Similarly, landmark identification on cephalometric radiographs is a typical task where AI can assist dentists (Schwendicke et al. 2021a; Schwendicke et al. 2019). Current AI-supported radiograph interpretation is as accurate or, in some cases, even more accurate than interpretation by practitioners. In addition, it saves time for the assessment and facilitates the generation of comprehensive and systematic reports.

Figure 2. AI and data-driven applications are situated all along the patient's dental journey. More details are provided in this section.
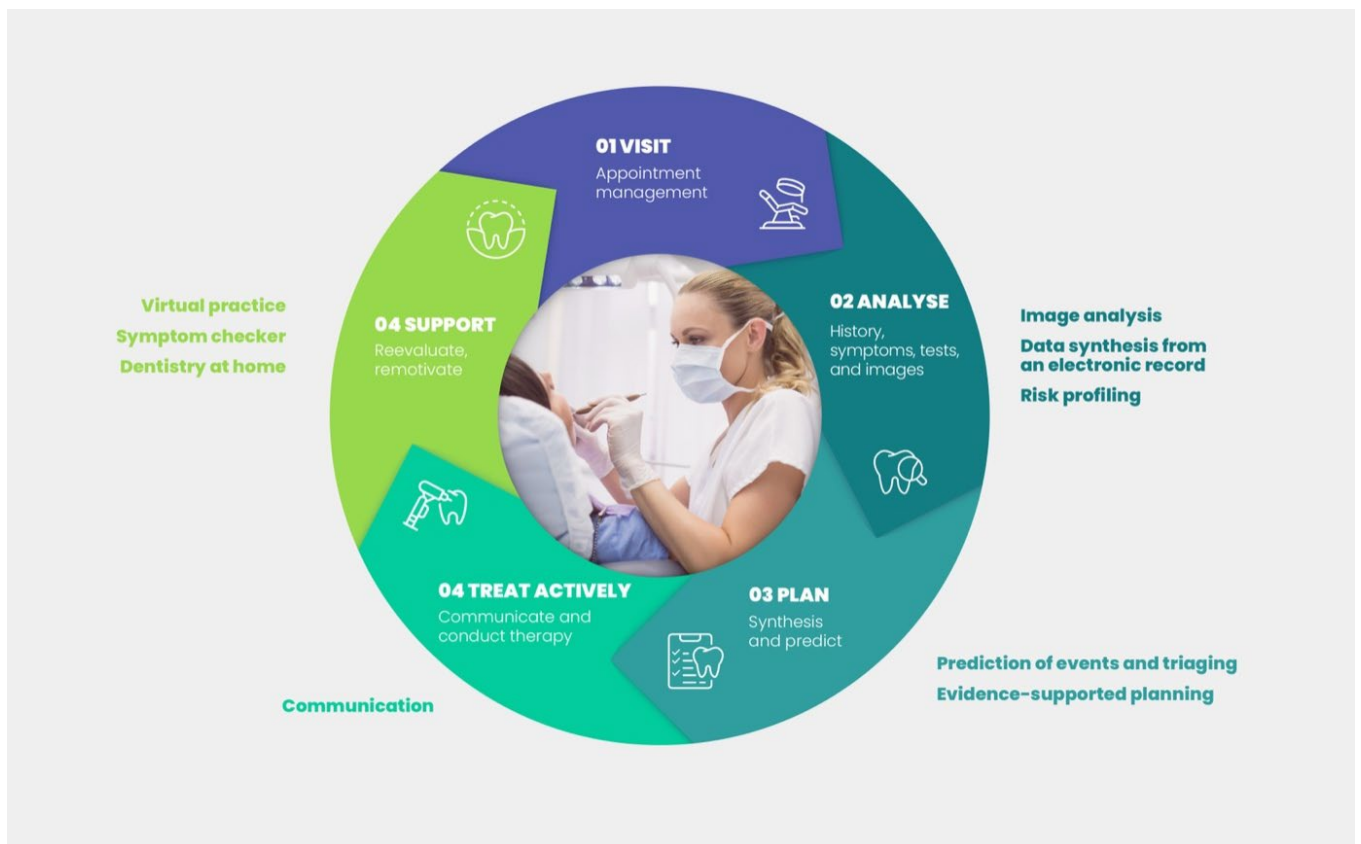
Several products are already on the market, while few have acquired the status of a medical product and are properly regulated. Dental professionals should critically evaluate the certification status of AI products before employing them for patient care. Standardization efforts are also most mature in this specific segment of AI applications as many organizations have provided or are preparing norms and guidance documents.

Similarly, the analysis of images has been widely attempted for non-radiographic data, such as 3D point cloud data from intraoral scans, photographs or near infrared transillumination images, to support smile design, orthodontic planning, preparation margin detection, restoration design or the detection of pathologies, for example caries or oral mucosa lesions.

Overall, future AI-based image analysis is likely to yield superior accuracy at higher speed and reliability than particularly unexperienced practitioners, supporting their decision-making process and allowing more standardized high-quality care. Moreover, AI will be able to guide practitioners' attention towards particularly difficult areas of an image, allowing for dedicated time for more in-depth assessment of these areas, and supporting them in communicating their findings with patients, for example via overlaying the affected areas with coloured pixel blobs (Kosan et al. 2022).

Data synthesis and Prediction

Clinical dentistry generates large amounts of data every day. Besides imagery, clinical data, historical data, claims, treatment data, and further data, for example from other diagnostic tests, are available. Given that many patients attend dentists regularly and repeatedly, there is a wealth of longitudinal data on a broad part of the population available, at least in many high-income countries. Currently, these data are oftentimes compiled in isolated data silos, remaining hidden from combined use, and are

difficult to be assessed jointly when evaluating the signs and symptoms of a patient or risks for future oral and dental conditions.

Making use of these data and leveraging them for better understanding of the patient, his or her risk profile and needs, is another promise of AI. Modern approaches to access the available data and merge and synthesize them, for example in dashboards, aim to provide the practitioner with the available data in more comprehensive and useful ways. Allowing a more holistic view of each patient with his or her individual history and risks is expected to improve the quality but also the efficiency of care, reducing the need for repeated and costly assessments.

Using data acquired over several visits and years, like imagery, clinical assessments, historical or data on medications or systemic conditions, will help to overcome the current "one size fits all approach" towards managing, diagnosing, and treating patients and assist moving to a more personalized, precise, and preventive care. Also, data provided by patients, for example dietary or toothbrushing behaviour collected by apps and mobile devices, will be increasingly employed, and patients will generally be more involved in the care process (as data donor, but also via apps and continuous support in the "virtual" practice). Hence, AI will support a more participatory dental care, too. The result is called P4 dentistry; a more personalized, precise, preventive and participatory dentistry (Flores et al. 2013).

Another unexplored data accumulation is implantable and wearable sensors. With the evolving biomedical sensor related technology, miniature and nano-sized sensors could be used, for example, to monitor saliva and, indirectly, oral health or disease. Moreover, saliva diagnostics and the related data could also be employed to assess other human diseases given the close link between saliva and overall health (Cui et al. 2022).

Notably, current AI towards risk profiling, predicting the future risk of caries, periodontitis,

progression of lesions, or tooth loss, is in its infancy. Understanding associations within larger datasets is often possible with AI, allowing identification of risk factors or indicators. Predicting risks, however, is far more difficult. In many cases, today's AI remains limited in its predictive power and generalizability, as outlined below (Schwendicke and Krois 2022).

Evidence-supported planning and conduct of therapies

As AI allows synthesizing data and identifying risk factors and patterns, it will further support evidence-based treatment planning and conduct. Providing more effective and efficient, individualized and person-centred care will not only facilitate consideration of more patient-specific aspects, but also incorporate external evidence-based information, for example from guidelines and standards of care. This will allow clinicians to standardize and calibrate treatment plans and foster reliable high-quality care. Recently, it has been highlighted that the clinical benefit of an AI, in this case a diagnostic support AI, would even benefit from such treatment recommendations, curbing overtreatment and allowing early intervention where needed (Mertens et al. 2021). As such, AI can improve the cost-effectiveness of oral health care (Schwendicke et al. 2020).

An additional benefit of more data-driven care is the option to regularly and objectively assess treatment needs, treatments provided, and outcomes yielded. Ultimately, this will foster value-based, but also more targeted care (for example when referral pathways need to be tailored for individual patients). Notably, it will mean that the quality of dental work will be objectifiable; insurance companies have already started to employ AI to detect subpar treatment and fraudulent insurance claims.

Patient Interaction

Given that the linkage of data is a major aim of data-driven care and AI tools for health, a more continuous interaction with patients is also likely to be facilitated. As many dental conditions are behaviourally grounded, such a more intimate interaction instead of what is often referred to as the "one-stop" approach in dental care may allow addressing dental conditions where they are rooted: in people's daily lives. A number of possible examples come to mind:

- Evaluating the toothbrushing behaviour of patients via their electronic toothbrush or assessing their dietary data via nutrition apps.

- Using symptom checkers or remote monitoring during orthodontic tooth movements.

- Providing the option for synchronous communication via mobile devices with a dental care professional or asynchronous communication via symptom checker apps or chatbots

In the future, one can expect patients to be more informed when they enter the practice. Conversely, dentists will know more about their patients' routines and the related risks. Moreover, many dental services, mainly around advice and supportive care, will be translated into the virtual room, increasing the efficiency of care but also its sustainability as the main driver of the dental carbon footprint is patients' travel to and from the dental office (Duane et al. 2014). Continuous monitoring will also allow regular feedback to facilitate quality improvements, patient engagement and empowerment, which are the bases for health and behaviour interventions. Lastly, AI will not only link patients and providers but also allow closer interaction between different providers, fostering intra- and interprofessional integrated care.

Section 3. Community and Public Health

Besides being used by dentists to manage the health of individual patients, AI and data-driven tools will enrich and help to expand community and public dental health services. For example, AI-based applications will allow closer linkage of outreach community-based services with the practice setting using remote synchronous and asynchronous communication and data exchange. This will allow expansion of the reach of traditional dental services to those with only limited or no access at present, increasing affordability, accessibility, and equity of care.

Moreover, data-driven technologies including AI will facilitate better information and decision-making on the dental public health level: Using standardized health indicators in oral health surveys as well as a more systematic and comprehensive reporting of the specific services provided and the resulting outcomes yielded within different health care systems or units will allow a more detailed picture of the public's dental health, the associated needs and demands, and the impact of services in terms of effectiveness and efficiency (Riordain et al. 2021). Comprehensive surveillance of health will build on regular and likely automated data generation. AI will allow leveraging the yielded large data pools for predicting health and demands guiding targeted workforce development (below), informed setup of services, benchmarking of healthcare interventions and policy, and overall reducing wasteful processes and ensuring sustainability of oral care.

Lastly, employing mobile tools powered by AI will enable more effective, individualized public health communication, facilitating the general public's and patients' involvement in oral and dental health and fostering better health for all.

Section 4. Education and the workforce

A major obstacle for delivering optimal care to all in the future will be workforce shortages and a maldistribution of the workforce. This will be true for high-income countries where ageing of the workforce comes in parallel with the feminization of it, leading to mass-retirement of dentists and reduced net availability for those who did not retire, for a population with sustained oral healthcare needs. However, it also applies to low- and middle-income countries where a fast-growing population's oral health care needs will exceed the workforce capacity but also the capability to train workforce fast enough. The oral and dental workforce is a key challenge identified in FDI's Vision 2030.

AI Impacts

AI can have an impact on dental education, expanding the workforce's capacities and scope of individual providers, but also workforce planning and allocation:

(1) AI will support a needs-based dynamic and adaptive workforce planning process. Integrating and extrapolating data from various sources such as oral health surveys, publicly available population, geostatistical and social data using AI will allow to capture age-group and spatially specific oral health, its translation into needs and demands and the required resources over time. It will allow workforce planning according to the population, its health and needs, at small spatial scale and in actionable units.

(2) AI tools will further enable the delivery of high-quality services by a diversified workforce, including dentists, other oral healthcare professionals or even providers from other professions or laypeople. AI-based support systems will facilitate more modularized models of care, with providers of different levels of competency jointly and synergistically providing care on-premises, but also at home or remotely.

Such models will provide physically and economically accessible basic services to all but also joint decision-making for cases of high complexity.

(3) Moreover, AI can be expected to have an impact on dental education methods and scope. Non-synchronous learning models, a more learner-centric provision of educational content as well as education being increasingly based on simulation including augmented or virtual reality-based teaching and training are promises of AI for dental education (Joda et al. 2019).

A major challenge to be highlighted is that the current dental workforce shows only limited digital literacy. Admission to dental training programs requires competencies in biological sciences, clinical and technical procedures, and the training itself also focuses on these aspects. Mathematical and informatics knowledge and skills are rarely required or reflected. Choosing between different AI tools for a specific practice, commissioning AI for community and public health purposes, and appraising AI with regards to its capacity and quality will require knowledge about the way AI works. Since data is an indispensable component of AI, the pitfalls around datasets underlying an AI application, but also around measuring the performance of an AI application must be known. Moreover, as laid out below, research around the impact of AI on dental providers, patients, and payers, is needed to better understand possible risks, for example of automation and complacency (Parasuraman and Manzey 2010), and how to mitigate them.

Section 5. Research

AI is already widely used in oral and dental research. The capabilities of AI to make use of complex data have facilitated a rapid uptake of various applications in the field, mainly for image analysis using neural networks (deep learning), but also for numerical data using a range of shallower machine learning approaches. Employing AI allows to identify patterns in large

omics datasets, but also routine or claims data with unprecedented accuracy. Notably, though, models from the AI realm do not consistently outperform less complex model approaches such as regression modelling, while the latter comes with lower computational demands and high explainability. Especially on oral and dental datasets of limited size, applying complex modelling may not be needed and may even come with limitations on its own.

Hence, to employ AI for oral and dental research most effectively, breaking up dental data silos and pooling data from multiple sites, levels, and modalities is required. Data pooling has been successfully employed in recent research activities around COVID-19 (Dayan et al. 2021), and allows development of AI models with high accuracy and, even more relevant, generalizability. As data protection concerns and regulations oftentimes hamper direct pooling, approaches like federated learning (where not data, but model parameters are pooled) may be needed (Bonawitz et al. 2019).

With the advent of the AI era in oral and dental research comes the need to and difficulties of appraising AI research. Many studies in the field suffer from poor reporting and low reproducibility. Even if sufficient detail is provided to reconstruct the applied methods, data and code as well as the computing environment are not regularly available, making replication impossible. The oral and dental research community is tasked with addressing these aspects. Researchers should adopt existing checklists and reporting standards for AI research, and reviewers and editors should contrast any submissions against these baseline reporting expectations (Schwendicke et al. 2021b). Moreover, peer-reviewing AI research comes with specific challenges that need to be addressed to ensure that the review process is as effective for oral and dental AI research as it is in other fields (Schwendicke et al. 2022). Lastly, the oral and dental research community is required to increase its digital literacy. Researchers but also the recipients of research,

dentists and other decision makers, need to be able to appraise AI research and the developed applications critically, as outlined.

Section 6. Risks and limitations of AI

The emergence and development of new technologies bring benefits and risks. For example, the development of the Internet has improved communication but problems such as infodemics have arisen that were difficult to foresee. The development of AI that allows machines to make automated decisions that humans previously made can significantly impact the way healthcare is delivered to people. While AI in dentistry could reduce human subjectivity in decisions and improve health, AI may contain biases that result in incorrect or discriminatory decisions in some population subgroups. The risks and limitations of AI may arise during the conceptualization stage of the problem to be solved with AI, during the development, implementation or finally during the evaluation of the AI model.

The fact that a specific technology exists does not mean that it should be applied to any problem. During the conceptualization of the problem to be solved by AI, it is required to first answer whether the new tool is necessary and what health problem it is intended to solve. To minimize the risks of generating superfluous tools for problems that already have a solution or unethical ones to increase an existing health inequity, it is best to involve all stakeholders, users, recipients, and payers of AI, in this phase, with a particular emphasis on minority or marginalized groups. Moreover, a comprehensive assessment of the available evidence towards existing standard of care solutions, their effectiveness, efficiency, and safety is needed. Research on AI in healthcare has in parts been decoupled from real-world problem solving, focusing on metrics and not clinical impact. The majority of AI for health interventions developed to date lack demonstrations of clinical or further benefit compared with the standard of care (Wynants et al. 2020; Zhou et al. 2021). Notably, this is also true for other innovations in healthcare.

The development of AI for health applications centres around the data used for training the AI. Often, small, unbalanced or unrepresentative datasets are employed, resulting in AI suffering from limited accuracy and generalizability as well as biases. This is rooted in data being siloed, limitedly interoperable, not standardized, and usually not available in sufficient breadth to yield representativeness. Recent calls for scientific data management demand data to be detectable, accessible, interoperable, and reusable by machines (FAIR principles) (Wilkinson et al. 2016). Several challenges during the development of AI for health applications arise:

1. Bias and limited generalizability

As outlined, many datasets used to train oral and dental AI lack inclusiveness and representativeness, including systematic under-representation of a specific gender, age group, race and others. This often results in AI replicating societal bias, for example discriminating under-represented and marginalized groups. Data may also suffer from biases introduced during the data collection, for example data collected by an insurer may reflect insurance priorities and be unsuitable data for developing and training AI for clinical purposes. Biases may also arise from data being collected from only a few populations. For example, in the USA, the majority (71%) of AI for health applications were trained on data from three states: California, Massachusetts, and New York (Kaushal et al. 2020). Data generated in clinical routine may itself suffer from confounding bias. In a recent study on radiographs to classify the severity of COVID-19 from chest radiographs, standing and lying patients were included. Because lying patients by large suffered from

more severe disease, the model erroneously learned to classify severity according to the patients' position (Roberts et al. 2021).

2. Accessibility

Biases and limited generalizability are often grounded in the overall limited and selective accessibility of data. In many cases, data is not available due to data protection reasons. A range of approaches to broaden the access to health data, improving AI performance, fairness, and generalizability, have been developed, for example broad consent, opt-in for data sharing rather than opt-out, or the conversion of health data into non-fungible tokens, which would allow patients to have control over who can access this data and under what circumstances. A technical approach to pool data siloed in different data lakes without sharing the data is federated learning (Bonawitz et al. 2019).

3. Interoperability

Even if data are accessible for developing and training an AI, limited interoperability of healthcare data remains a significant hurdle. The two major aspects to consider are syntactic and semantic interoperability. Syntactic interoperability defines the format and structure of data. It is backed by the International Standards Development Organizations such as Health Level Seven International (HL7) and its Fast Health Interoperability Resource (FHIR). Semantic interoperability is ensured by agreed terminologies, nomenclatures, and ontologies being employed, such as SNOMED CT, the Logical Observation Identifiers Names and Codes (LOINC), or the identification of Medicinal Products for medicines (IDMP). In oral and dental research, none of these concepts are particularly popular or widespread. Even terminologies and classifications defined by dental experts themselves are not necessarily adopted or interchangeably usable. An example is differing the World Health Organization (WHO) and the International Caries Detection and

Assessment System (ICDAS) criteria for caries diagnosis. For omics data, the creation of a global database repository with interoperable microbiome sequences and annotated host metadata would be similarly desired (Marcos-Zambrano et al. 2021). Interoperability of oral data and its collection in routine settings (e.g. electronic health records) but also prospective studies (e.g. national oral surveys) should be aimed for, and agreed standards of captured data (for example outcomes in interventional studies) may assist (Riordain et al. 2021).

4. Truthfulness

The validity and quality of the data but also the labelling (annotations) must be carefully considered. Often, a hard gold standard such as histological assessment is not available. Instead, multiple human annotators provide the reference test (label), for example towards the presence of pathology on an image. The qualification of these annotators is critical to the eventual quality of any AI as the AI will not supersede the joint intelligence of the annotators. It remains unclear how disagreements between experts are best mitigated. Current concepts include consensus discussion and hierarchical reviewing.

In the implementation and maintenance stage, clinicians and healthcare staff play a critical role as any AI is no longer appraised as a model but as a healthcare intervention. Here, aspects around usefulness, acceptability, implementation, and maintenance of the resulting software product are relevant. Evaluation of the AI for its usefulness is critical at this stage. For example, during the COVID-19 pandemic, hundreds of models were developed to predict the risk for COVID-19 but none have been found clinically useful by a recent systematic review (Roberts et al. 2021). During the evaluation of an AI application, a wider set of performance metrics generated in clinical care should be considered, and any reporting of these metrics should follow established standards. Notably, the metrics should be clinically relevant

for decision-making, and the evaluation should reflect the interaction between the user and the AI software. Oftentimes, this interaction is a source of performance decreases compared with the performance determined by the AI itself.

Implementation and Maintenance

In the implementation and maintenance stage, clinicians and healthcare staff play a critical role as any AI is no longer appraised as a model but as a healthcare intervention. Here, aspects around usefulness, acceptability, implementation, and maintenance of the resulting software product are relevant. Evaluation of the AI for its usefulness is critical at this stage. For example, during the COVID-19 pandemic, hundreds of models were developed to predict the risk for COVID-19 but none have been found clinically useful by a recent systematic review (Roberts et al. 2021). During the evaluation of an AI application, a wider set of performance metrics generated in clinical care should be considered, and any reporting of these metrics should follow established standards. Notably, the metrics should be clinically relevant for decision-making, and the evaluation should reflect the interaction between the user and the AI software. Oftentimes, this interaction is a source of performance decreases compared with the performance determined by the AI itself. The generalizability of the AI should be displayed, and the underlying logic of the AI to come to a decision should be critically evaluated. Many AI applications are inherent black boxes. "Unboxing" them (making them explainable) and allowing clinicians to scrutinize the criteria the AI used for its decision is relevant to increase trust, but also to identify biases ("explainability").

Further aspects to consider are the derived long-term consequences of using an AI application and its impact on the user, the recipient, treatment pathways, and cost-effectiveness. Misuse of AI applications beyond the intended use should further be cautioned against.

Currently, data on these aspects are extremely limited or fully absent in the oral and dental arena. Given that technology has been identified as major driver of increasing healthcare cost (Nghiem and Connelly 2017), especially economic analyses are warranted (Chandra and Skinner 2012).

A summary of the risks and challenges is provided in the table 1.

Table 1. Risks at different stages of AI application development for dentistry.

Stage	Examples	Potential solution
Conceptualization	• Non-ethical application (e.g., unfair, increasing inequity)	Involve ethics experts in the conception of AI, consider representativeness of datasets to be used for training.
	• Application without scientifically grounded and medically useful objective (e.g., solely to increase revenue at patients' health expense; AI solution where the standard of care works close to perfect)	Involve patients and minority groups in the conceptualization and development of the model. Critically assess the objective of the AI for patients, providers, healthcare system.
Development	• Usage of biased, inaccurate, or non-representative data	Define minimal data selection criteria and evaluate data quality.
	• Limited data availability, accessibility and linkability	Aspire to an open research culture in which data are shared as openly as legal and ethical obligations permit. Adherence to FAIR principles, to make data detectable, accessible, interoperable and reusable by humans and machines; enforce interoperability (semantical and syntactical). Employ federated learning approaches. Creation of curated, high quality, heterogenous benchmarking datasets.
Implementation and maintenance	• AI performance is unexpectedly low.	Assess a comprehensive set of core metrics before deployment. Follow checklists for reporting of development and outcomes of AI in oral and dental research. Allow critical appraisal by third parties, e.g. during peer-review, by providing input data, insights on operational framework, or options to run the model (e.g. API, minimal viable product).
	• AI does have positive diagnostic effect, but not health effect.	Evaluate performance in different populations and population subgroups to gauge generalizability; continue to observe performance during usage (post market surveillance). Consider wider effects beyond diagnostic performance, i.e., the effect on care processes, patient-provider interaction, effects on treatment decisions and outcomes, cost-effectiveness.
	• AI performance is high in test dataset, but unexpectedly low in external datasets.	Employ explainable AI tools to assess AI reasoning and identify potential biases.
	• AI is not accepted by users, not properly implemented or maintained long-term.	Develop a core curriculum for oral dental staff to be able to critically assess AI tools and appraise their quality, but also to highlight risk of complacency bias, for example.
	• AI is misused and hence does not generate expected health or process gains.	Have users trained, safeguard AI by incorporating mitigation strategies against misuse, communicate uncertainty of AI outcomes to users.

Section 7. Governance

While there are opportunities and potential benefits associated with the use of AI in healthcare, there are the outlined risks (World Health Organization (WHO) 2021).

WHO's Consensus Principles on AI

The World Health Organization guidance provides six consensus principles to ensure that AI works for people's benefit:

1. Protect human autonomy: The use of AI or other computer systems should not diminish human autonomy, so humans should remain in control of the healthcare system and medical decisions.

2. AI must not harm people: Regulatory requirements and, generally, governance are needed to ensure the safety, accuracy, and efficacy of AI for health.

3. Ensuring transparency, explainability, and intelligibility: AI-based technologies must be understandable to developers, medical professionals, patients, users, and regulators. AI must be sufficiently documented prior to clinical implementation, and all methods or processes should allow humans to understand and trust the results.

4. Fostering responsibility and accountability: Humans require clear, transparent specifications of the tasks that an AI can perform and the conditions under which the AI can achieve this task. Humans should supervise the AI upstream, and downstream accountability remains with the human user.

5. Ensuring inclusiveness and equity: AI for health should ensure the widest possible appropriate, equitable use and access, irrespective of age, sex, gender, income, race, ethnicity, sexual orientation, ability, or other characteristics protected under human rights codes. This is achieved by making AI-based technologies available to the most disadvantaged and marginalized groups. Also, the inevitable power disparities that emerge including providers, patients, policymakers, governments and companies that occupy, develop, and implement AI-based health technologies must be minimized by assessing and monitoring any disproportionate effects that may maintain or increase a disproportionate share of power between providers, patients, policymakers, governments, and companies.

6. Promoting responsive and sustainable AI: Responsiveness requires adequate continuous, systematic, and transparent monitoring of AI technologies during their real-life use. The environmental impact of AI should be carefully and comprehensively assessed, including the data generation, training, implementation, inference, and maintenance associated with the AI application, and should be transparently communicated to the involved stakeholders. Sustainability also requires governments and companies to anticipate and address the disruptions caused by the implementation of AI technologies, for example their impact on users and recipients, the healthcare process, or wider policies needed to be implemented alongside an AI.

Implemented Governance Strategies

Reflecting these principles, various governance strategies have been developed and implemented globally. For instance, the EU Commission carved out seven requirements towards AI applications in general: (1) Human agency and oversight, including fundamental rights, human agency and human oversight; (2) Technical robustness and safety, including resilience to attack and security, fall back plan and general safety, accuracy, reliability, and reproducibility; (3) Privacy and data governance, including respect for privacy, quality and integrity of data, and access to data; (4) Transparency, including traceability, explainability and communication; (5) Diversity, non-

discrimination, and fairness, including the avoidance of unfair bias, accessibility and universal design, and stakeholder participation; (6) Societal and environmental well-being, including sustainability, and environmental friendliness, social impact, society, and democracy; (7) Accountability, including auditability, minimization, and reporting of negative impact, trade-offs and redress. All requirements are equally important, support each other, and should be implemented and evaluated throughout the AI system`s lifecycle. For regulation, the EU lawmakers set criteria differentiating "low risk" from "high risk" AI applications. They propose "a risk-based approach, according to which a given AI application should be considered high risk if both the target sector (e.g. healthcare) and the intended use involve significant risk (e.g. injury, death)". It was proposed "that operators of a high-risk AI system be subject to strict liability for any harm or damage caused by a physical or virtual activity, device or process driven by that AI system and be subject to a mandatory insurance regime…AI systems not listed as a high-risk AI system would, in principle, remain subject to fault-based liability, unless stricter national laws and consumer protection legislation is in force…"(4). The EU "Medical Devices Regulation" (MDR) operationalizes this risk-based approach, with medical devices being divided into different risk classes. Depending on the risk class of the product, a different conformity assessment procedure is foreseen before the product can be placed on the EU market. Low-risk medical devices are not subject to a pre-market authorization by a regulatory authority. Medium- and high-risk devices require a conformity assessment procedure, involving an independent third party as "notified body".

In the United States, two different authorities regulate AI for health technologies; the Centers for Disease Control and Prevention (CDC) and the Food and Drug Administration (FDA). Both are supervised by the US Department for Health and Human Services (HHS). HSS issued the Health Insurance Portability and Accountability

Act (HIPAA) in 1996, a federal law "that required the creation of national standards to protect sensitive patient health information from being disclosed without the patient's consent or knowledge". HHS issued the HIPAA Privacy Rule to implement the requirements. The Privacy Rule standards address the use and disclosure of individuals' health information and is regulated by the CDC. The FDA is responsible, among others, "for protecting the public health by ensuring the safety, efficacy, and security of human and veterinary drugs, biological products, and medical devices". Beyond that the FDA is "responsible for advancing the public health by helping to speed innovations that make medical products more effective, safer, and more affordable". In 2019, the FDA proposed a regulatory framework for modifications to AI/ML-based "Software as a Medical Device (SaMD)", and in January 2021, the FDA published a "SaMD-Action Plan" that includes; a proposed regulatory framework, a harmonized development of "Good Machine Learning Practice (GMLP)", a patient-centred approach on the role of transparency to users of AI/ML devices, a support of regulatory science efforts to develop a methodology for evaluation and improvement of ML algorithms concerning robustness and resilience, and an advance of real-world performance pilots. FDA's vision is that, with appropriately tailored total product lifecycle-based regulatory oversight, AI/ML-based SaMD will deliver safe and effective functionality that improves the quality of care. FDA described a "Predetermined Change Control Plan" in premarket submissions. In this approach, FDA expressed an expectation for transparency and real-world performance monitoring by manufacturers that could enable FDA and manufacturers to evaluate and monitor a software product from its pre-market development through post-market performance. With "Good Machine Learning Practice (GMLP)" the FDA describes a set of AI/ML-based best practices - data management, feature extraction, training, interpretability, evaluation, and documentation – that are akin to good software

engineering practices or quality system practices.

Conclusions for clinical practice

Overall, AI has implications for clinical practice today, but more so in the future:

1. Current applications focus on diagnostic support, mainly image analysis. Users should critically appraise the accuracy of diagnostic support systems, the data underlying the trained model and its testing as well as the resulting generalizability. Explainability of any application should be demanded, as the final responsibility for any decisions emanating from using a diagnostic support system remains with the user, i.e. the oral healthcare provider.

2. Emerging applications which may enter the practice soon employ speech analysis (natural language processing), facilitating speed-based reporting in clinical care, but also allowing rapid and comprehensive mining of electronic health records etc. Such analyses may form the backbone of health information systems (dashboards) and support integrated care models.

3. Personalized, precise, preventive, and participatory care (P4 dentistry) (Hood and Flores 2012) will need to employ data from a range of sources and relies heavily on predictive modelling. Current predictive models are of limited accuracy and generalizability, oftentimes as the datasets used for training them are small and stem from one or few sources. P4 dentistry remains a promise of the future for now.

Call for action

Based on the outlined four application areas and the associated challenges, the oral health community will need to actively engage in a number of areas:

• To make sure that AI applications acquire maximal usefulness and generalizability, action is needed to improve access to data and foster its usage for training and testing. Data protection concerns and reservations against the commercial use of healthcare data need to be carefully balanced against the potential benefit of AI for patients and the society. Technological approaches like federated learning should be actively applied to dental AI tasks, and harmonization of data to improve interoperability should be actively pursued.

• AI applications with demonstrated benefit for patients, providers or the healthcare system should be actively promoted, and its uptake incentivized. Notably, the impact of these applications on all stakeholders should be critically appraised.

• Generally, the evidence supporting AI applications for oral and dental care needs strengthening. The active development of AI should be complemented by research into the true usefulness, applicability and implementation of these applications.

• This evidence needs to be utilized by the community; providers and healthcare decision makers need to acquire a basic literacy around data, digital applications and AI. For providers, the decision to use or not use an AI application should be taken on an informed basis – similar to the decision to use or not use an adhesive system, for example. An international core curriculum should be developed and adapted to national needs, and educators should be actively trained to teach such curriculum in graduate or postgraduate training.

The community should engage into the widespread activities around standardization. The development of standards towards AI will serve as safeguards, and active engagement of oral and dental stakeholders is indispensable if these standards are expected to serve the oral and dental health community.

References

1. Bell G, Hey T, Szalay A. 2009. Beyond the data deluge. Science. 323(5919):1297.

2. Bonawitz K, Eichner H, Grieskamp W, Huba D, Ingerman A, Ivanov V, Kiddon C, Konečný J, Mazzocchi S, McMahan HB. 2019. Towards federated learning at scale: System design. arXiv preprint arXiv:190201046.

3. Chandra A, Skinner J. 2012. Technology growth and expenditure growth in health care. Journal of Economic Literature. 50(3):645-680.

4. Cui Y, Yang M, Zhu J, Zhang H, Duan Z, Wang S, Liao Z, Liu W. 2022. Developments in diagnostic applications of saliva in human organ diseases. Medicine in Novel Technology and Devices. 13:100115.

5. Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, Liu A, Costa AB, Wood BJ, Tsai C-S et al. 2021. Federated learning for predicting clinical outcomes in patients with covid-19. Nature Medicine. 27(10):1735-1743.

6. Duane B, Taylor T, Stahl-Timmins W, Hyland J, Mackie P, Pollard A. 2014. Carbon mitigation, patient choice and cost reduction--triple bottom line optimisation for health care planning. Public Health. 128(10):920-924.

7. Flores M, Glusman G, Brogaard K, Price ND, Hood L. 2013. P4 medicine: How systems medicine will transform the healthcare sector and society. Personalized medicine. 10(6):565-576.

8. Hood L, Flores M. 2012. A personal view on systems medicine and the emergence of proactive p4 medicine: Predictive, preventive, personalized and participatory. N Biotechnol. 29(6):613-624.

9. Joda T, Gallucci GO, Wismeijer D, Zitzmann NU. 2019. Augmented and virtual reality in dental medicine: A systematic review. Comput Biol Med. 108:93-100.

10. Kaushal A, Altman R, Langlotz C. 2020. Geographic distribution of us cohorts used to train deep learning algorithms. Jama. 324(12):1212-1213.

11. Kosan E, Krois J, Wingenfeld K, Deuter CE, Gaudin R, Schwendicke F. 2022. Patients' perspectives on artificial intelligence in dentistry: A controlled study. J Clin Med. 11(8).

12. LeCun Y, Bengio Y, Hinton G. 2015. Deep learning. Nature. 521(7553):436-444.

13. Marcos-Zambrano LJ, Karaduzovic-Hadziabdic K, Loncar Turukalo T, Przymus P, Trajkovik V, Aasmets O, Berland M, Gruca A, Hasic J, Hron K et al. 2021. Applications of machine learning in human microbiome studies: A review on feature selection, biomarker identification, disease prediction and treatment. Frontiers in Microbiology. 12.

14. Mertens S, Krois J, Cantu AG, Arsiwala LT, Schwendicke F. 2021. Artificial intelligence for caries detection: Randomized trial. Journal of dentistry. 115:103849.

15. Nghiem SH, Connelly LB. 2017. Convergence and determinants of health expenditures in oecd countries. Health Econ Rev. 7(1):29.

16. Dictionaries. 2022. https://www.oxfordlearnersdictionaries.com/definition/english/artificial-intelligence; [accessed].

17. Parasuraman R, Manzey DH. 2010. Complacency and bias in human use of automation: An attentional integration. Hum Factors. 52(3):381-410.

18. Riordain RN, Glick M, Mashhadani S, Aravamudhan K, Barrow J, Cole D, Crall JJ, Gallagher JE, Gibson J, Hegde S et al. 2021. Developing a standard set of patient-centred outcomes for adult oral health - an international, cross-disciplinary consensus. Int Dent J. 71(1):40-52.

19. Roberts M, Driggs D, Thorpe M, Gilbey J, Yeung M, Ursprung S, Aviles-Rivero AI, Etmann C, McCague C, Beer L et al. 2021.

Common pitfalls and recommendations for using machine learning to detect and prognosticate for covid-19 using chest radiographs and ct scans. Nature Machine Intelligence. 3(3):199-217.

20. Schwendicke F, Chaurasia A, Arsiwala L, Lee JH, Elhennawy K, Jost-Brinkmann PG, Demarco F, Krois J. 2021a. Deep learning for cephalometric landmark detection: Systematic review and meta-analysis. Clin Oral Investig.

21. Schwendicke F, Golla T, Dreher M, Krois J. 2019. Convolutional neural networks for dental image diagnostics: A scoping review. Journal of Dentistry.103226.

22. Schwendicke F, Krois J. 2021. Data dentistry: How data are changing clinical care and research. J Dent Res.220345211020265.

23. Schwendicke F, Krois J. 2022. Precision dentistry-what it is, where it fails (yet), and how to get there. Clin Oral Investig. 26(4):3395-3403.

24. Schwendicke F, Marazita ML, Jakubovics NS, Krois J. 2022. Big data and complex data analytics: Breaking peer review? Journal of dental research.220345211070983.

25. Schwendicke F, Rossi JG, Göstemeyer G, Elhennawy K, Cantu AG, Gaudin R, Chaurasia A, Gehrung S, Krois J. 2020. Cost-effectiveness of artificial intelligence for proximal caries detection. Journal of dental research.22034520972335.

26. Schwendicke F, Singh T, Lee JH, Gaudin R, Chaurasia A, Wiegand T, Uribe S, Krois J. 2021b. Artificial intelligence in dental research: Checklist for authors, reviewers, readers. J Dent.103610.

27. Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, Blomberg N, Boiten J-W, da Silva Santos LB, Bourne PE et al. 2016. The fair guiding principles for scientific data management and stewardship. Scientific Data. 3:160018.

28. Ethics and governance of artificial intelligence for health. 2021. https://www.who.int/publications/i/item/97892 40029200; [accessed].

29. Wynants L, Van Calster B, Collins GS, Riley RD, Heinze G, Schuit E, Bonten MMJ, Dahly DL, Damen JA, Debray TPA et al. 2020. Prediction models for diagnosis and prognosis of covid-19: Systematic review and critical appraisal. BMJ. 369:m1328.

30. Zhou Q, Chen Z-h, Cao Y-h, Peng S. 2021. Clinical impact and quality of randomized controlled trials involving interventions evaluating artificial intelligence prediction tools: A systematic review. npj Digital Medicine. 4(1):154.